

HEPiX Autumn 2023 Workshop

Monday, 16 October 2023 - Friday, 20 October 2023



Book of Abstracts

Contents

Diamond Light Source Site Report	1
Migrating to Slurm from Grid Engine: Politics, Partitions and Problems	1
Graylog-as-a-service: using Kubernetes to rescue a service from ancient hardware	1
From Generative to Interactive AI: Towards Artificial General Intelligence? Use on Local Data and Applications Examples	2
MFA for SSH at Diamond Light Source	2
Ceph in 2023 and Beyond	3
Maintaining a legacy data centre	3
Data centre adventures during building renovations	3
AGLT2 Site Report	4
The Design and Progress of Data Management and Data Service for HEPS	4
USATLAS SWT2 Center Site Report	5
Update on ARM for WLCG	5
Carbon negative computing?	5
RAL Site Report	6
ASGC Site Report	6
Building a cloud-native ATLAS Tier 2 on Kubernetes	6
Our new router is a Nokia, but can it play snake?	7
KEK Site Report	7
Sustainable Self-Inspection Initiatives for Improving Information Security at KEK	7
HEPiX Benchmarking Working Group Report	7
The new BaBar Long Term Data Analysis facility	8
Configuration management in the PDP group at Nikhef	8
IHEP Site Report	9

Exploring storage technologies for HPSS disk caches	9
Canadian ATLAS Tier-1 Site Report	9
Canadian ATLAS Tier-1 Analytics Infrastructure	10
Deploying and Running Ceph Clusters for Analysis Facilities at RAL	10
Securing the RAL campus	10
Supporting Distributed Subatomic Physics Computing in Canada	11
Quantum Assisted Calorimeter Simulation	11
Sustainable Immersion Cooling for HPC: The Path to Energy Efficiency and Environmental Responsibility	12
Ensuring Use of IPv6 after it is deployed	13
New Security Trust and Policies - for WLCG and other Research Infrastructures	13
LHCOPN and LHCONE update	13
CERN site report	13
Oxford Computer Room Air Conditioning upgrades	14
BNL Site Report	14
The 2023 HTCondor workshop in Europe	14
Ocean Networks Canada: Continuously Delivering Multidisciplinary Data from the Deep	15
LHEP site report	15
The CERN IT Linux strategy (after the recent events in the EL eco-system)	15
The P-One ocean-based neutrino experiment - status and prospects	16
The Digital Humanities, and its Foundation for Open Social Scholarship	16
An Square Kilometer Array Regional Centre: Scaling Digital Research Infrastructure for Astronomy in Canada	16
Overview of the Coffea-Casa Analysis Facility hosted at the University of Nebraska-Lincoln	17
The Digital Research Alliance of Canada	17
InvenioRDM at the SDCC	17
The DESY User Registry	18
Zimbra at DESY – what comes next?	18
An even more Efficient Nordic Dcache Interface to TSM	18
FZU Site Report	18

WLCG Networking Activities, Status and Plans	19
Deploying dCache in Kubernetes	19
Nikhef overview and site report	19
HIP site report	20
Welcome	20
Introduction to UVic and Logistics	20
An introduction to HEPiX	20
Linux at DESY	20
Discussion	21
IDAF@DESY: Status & Outlook	21
DDN - Importance of checkpoints with LLM	21
Vendor Introduction - Weka	21
NET2: a first example of OpenShift/OKD for Tier 2 provisioning and cluster management in US ATLAS	21
Election of third co-chair	21
Closing Remarks	22
Call for participation	22

Site Reports 2 / 1**Diamond Light Source Site Report****Author:** James Thorne¹¹ *Diamond Light Source***Corresponding Author:** j.i.thorne@gmail.com

Latest news from Diamond Light Source. There are a number of updates since the last talk in Umeå, including:

- Migration to Slurm from Grid Engine
- Implementation of MFA for SSH and No Machine
- Graylog-as-a-service
- Diamond-II update
- A power outage and recovery

Computing and Batch Services / 2**Migrating to Slurm from Grid Engine: Politics, Partitions and Problems****Author:** Murray Collier^{None}**Co-author:** James Thorne¹¹ *Diamond Light Source***Corresponding Authors:** j.i.thorne@gmail.com, murraycollier@gmail.com

Diamond Light Source have migrated to Slurm from Univa Grid Engine this year. We will present a summary of our challenges and solutions including:

- Catering for multiple data centres and storage systems
- Ensuring stakeholder buy-in
- Supporting automated submission systems
- Accounting
- Auto-creation of user accounts
- Elasticsearch accounting
- Migrating to a new deployment and configuration system
- Node health checks

Basic IT Services and End User Services 2 / 3**Graylog-as-a-service: using Kubernetes to rescue a service from ancient hardware**

Author: Sonia Taneja¹

Co-author: James Thorne²

¹ *Diamond light source*

² *Diamond Light Source*

Corresponding Author: sonia91taneja@gmail.com

Diamond Light Source had a single, monolithic, ancient version of Graylog running on even more ancient hardware. We present the migration of this service to Graylog-as-a-Service: an instance per user community running on Kubernetes. Benefits include:

- User communities can manage their own instances of Graylog
- User communities can no longer deny service to each other by flooding Graylog.
- Some redundancy comes “for free” by using our Kubernetes infrastructure.
- We can use our Elasticsearch backend for log storage

Computing and Batch Services / 4

From Generative to Interactive AI: Towards Artificial General Intelligence? Use on Local Data and Applications Examples

Author: Imed Magroune¹

¹ *CEA/DRF/IRFU/DEDIP//LIS*

Corresponding Author: imed.magroune@cea.fr

Abstract:

In the rapidly evolving world of Artificial Intelligence (AI), Large Language Models (LLMs) have emerged as a powerful tool capable of understanding, interpreting, and generating human-like text. This presentation will delve into the intricacies of state-of-the-art models such as GPT, LLAMA, ALPACA and Orca, highlighting their unique capabilities and their potential in transforming High Energy Physics IT.

The talk will explore the practical aspects of fine-tuning these models on local servers using local data, addressing the technical challenges and considerations, and providing effective solutions. We will discuss the potential benefits and the flexibility that local fine-tuning brings to the table, especially for HEPiX, where data interpretation is of paramount importance.

Furthermore, the presentation will showcase real-world examples and case studies to illuminate the practical applications of these models in the HEPiX field. It aims to demonstrate how these cutting-edge AI models can be utilized to comprehend complex HEP data and generate meaningful insights.

This talk invites all HEPiX participants and stakeholders to consider the potential of LLMs as a robust tool for data interpretation and knowledge generation, and encourages a discussion on further exploration and collaboration in this exciting intersection of AI and High Energy Physics.

In this era where data is the new oil, let us tap into the potential of Large Language Models to refine our data and generate valuable insights for High Energy Physics.

Networking & Security 3 / 5

MFA for SSH at Diamond Light Source

Author: James Thorne¹

¹ *Diamond Light Source*

Corresponding Author: j.i.thorne@gmail.com

Diamond Light Source is implementing MFA for SSH and No Machine. This is a story of our trials with PAM, RADIUS, Microsoft and Google Authenticator. I'll present the solutions considered as well as the pros and cons of each, particularly with regard to the difficulties faced regarding MFA for facility users.

Storage and Filesystems / 6

Ceph in 2023 and Beyond

Author: Dan van der Ster¹

¹ *Clyso*

Corresponding Author: dan.vanderster@clyso.com

Ceph is a popular software defined storage system providing an open source alternative to proprietary appliances and cloud storage. It provides block and object storage for on-premises clouds as well as networked filesystems for shared compute facilities including several WLCG sites. Altogether, Ceph aims to be a single solution to all of our data centre storage needs – the “Linux of Storage”.

This talk will present the status of the open source Ceph project, recent improvements in the latest (Reef) release and outline the future vision for the project.

IT Facilities & Business Continuity / 7

Maintaining a legacy data centre

Author: Floris Bieshaar¹

¹ *Nikhef*

Corresponding Author: florisb@nikhef.nl

A brief rundown of the facilities/technical installation for our in-house data centre as well as our semi commercial Nikhef Housing data centre.

This talk should take approximately 25 minutes

IT Facilities & Business Continuity / 8

Data centre adventures during building renovations

Author: Floris Bieshaar¹

¹ *Nikhef*

Corresponding Author: florisb@nikhef.nl

A rundown on how we've maintained continuity during a full building renovation and data centre expansion.

Site Reports 2 / 9

AGLT2 Site Report

Author: Wenjing Dronen^{None}

Co-author: Shawn Mckee¹

¹ *University of Michigan*

Corresponding Author: wuwj@umich.edu

We will report the site's overall status and its recent activities, including setting up the SOC on EL9 and using Redhat Satellite server to provision for the RHEL 9 systems, and transitioning our software from CentOS 7 to RHEL9.

Storage and Filesystems / 10

The Design and Progress of Data Management and Data Service for HEPS

Author: Hao Hu¹

Co-authors: Bo Zhuang¹; Qi Fazhi²; qi luo³; 王浩帆 hfwang²

¹ *Institute of High Energy of Physics*

² *IHEP*

³ 中科院高能物理所计算中心

Corresponding Author: huhao@ihep.ac.cn

China's High Energy Photon Source (HEPS), the first national high-energy synchrotron radiation light source and soon one of the world's brightest fourth-generation synchrotron radiation facilities, is being under intense construction in Beijing's Huairou District, and will be completed in 2025. The 14 beamlines for the phase I of HEPS will produce more than 300PB/year raw data. Efficiently storing, analyzing, and sharing this huge amount of data presents a significant challenge for HEPS.

To make sure that the huge amount of data collected at HEPS is accurate, available and accessible, we developed an effective data management system that is aimed at automating the organization, transfer, storage, distribution and sharing of the data produced from HEPS experiments. First, the general situation of HEPS and the construction progress of the whole project are introduced. Second, the architecture and data flow of the HEPS DMS are described. Third, key techniques and new function modules implemented in this system are introduced. For example, the process of automatic data tracking when using a hierarchical storage policy is illustrated, and how the DMS deals with the metadata collection when an emergency occurs such as beamline network interruption. Finally, the progress and the effect of the data management and data service system deployed at testbed beamlines of BSRF are given.

The integration and the verification of the whole system at 3W1 beamline of BSRF (Beijing Synchrotron Radiation Facility) were finished and achieved great success. It strongly proved the rationality of the design scheme and the feasibility of the technologies. After the optimization and upgrade of the functionality, the data management system were deployed at 4W1B, which is a running beamline at BSRF, can provide data service for beamline users.

Site Reports 1 / 11

USATLAS SWT2 Center Site Report

Author: Horst Severini¹

¹ *University of Oklahoma (US)*

Corresponding Author: hs@nhn.ou.edu

I will give a status update of the USATLAS SWT2 Center.

Computing & Batch Services / 13

Update on ARM for WLCG

Author: David Britton¹

Co-authors: Bruno Borbely ¹; Dwayne Spiteri ²; Emanuele Simili ¹; Gordon Stewart ; Samuel Cadellin Skipsey

¹ *University of Glasgow (GB)*

² *University of Glasgow*

Corresponding Author: david.britton@cern.ch

Over the last 18 months we have investigated the use of the 80-core Ampere Altra for WLCG workloads and have previously reported that for WLCG workloads this ARM-based machine delivers significant energy-savings whilst being comparable in both speed and cost to typical AMD machines. More recently we have extended this work to the 128-core Ampere Altra Max and will present these results for the first time. In the meantime, the installation of a 2000-core Altra ARM-farm at Glasgow has allowed significant in-house ARM resources to be presented on the WLCG for the first time. The facility is being validated by ATLAS, re-running a recent google validation of their simulation workload and running reconstruction in tandem. Looking ahead, Glasgow is in the process of procuring a NVIDIA Grace (ARM-based) processor to characterise its performance (though we expect this to push performance rather than energy-efficiency) and hope to have new results to present on the AMD Bergamo processor available for the meeting.

IT Facilities & Business Continuity (C&F) / 14

Carbon negative computing?

Author: Erik Mattias Wadenstein¹

¹ *University of Umeå (SE)*

Corresponding Author: mattias.wadenstein@cern.ch

A few methods of calculating CO₂ emissions from computing, centred around the local circumstances at HPC2N, Umeå University. Can we actually have carbon negative scientific computing?

Site Reports 2 / 15

RAL Site Report

Author: Martin Bly¹

¹ *STFC-RAL*

Corresponding Author: martin.bly@stfc.ac.uk

An update on activities at RAL.

Site Reports 2 / 16

ASGC Site Report

Author: Eric Yen¹

Co-author: Felix Lee

¹ *ASGC*

Corresponding Author: eric@sinica.edu.tw

Updates of WLCG, user community collaborations and technical solutions etc. at ASGC.

Grid, Cloud & Virtualisation and Operating Systems / 17

Building a cloud-native ATLAS Tier 2 on Kubernetes

Author: Ryan Taylor¹

¹ *University of Victoria (CA)*

Corresponding Author: rptaylor@uvic.ca

The University of Victoria operates an Infrastructure-as-a-Service scientific cloud for Canadian researchers, and a Tier 2 WLCG site for the ATLAS experiment at CERN. Over time we have taken steps to migrate the Tier 2 grid services to the cloud. This process has been significantly facilitated by basing our approach on Kubernetes. We have exploited the batch capabilities of Kubernetes to run grid computing jobs and replace the conventional grid computing elements by interfacing with the Harvester workload management system of the ATLAS experiment. We have also adapted and migrated the APEL accounting service and Squid caching proxies to cloud-native deployments on Kubernetes, and are prototyping a Kubernetes-based grid storage element. We aim to enable fully comprehensive deployment of a complete ATLAS Tier 2 site on a Kubernetes cluster via Helm charts. We also describe our experience running a high-performance self-managed Kubernetes ATLAS Tier 2 cluster at the scale of 8,000 CPU cores for several years, and compare with the conventional setup of grid services.

Networking & Security 2 / 19**Our new router is a Nokia, but can it play snake?**

Author: Bart van der Wal¹

Co-author: Tristan Suerink²

¹ *Nikhef*

² *Nikhef National institute for subatomic physics (NL)*

Corresponding Author: bwal@nikhef.nl

An overview of why Nikhef decided to buy two 7750-SR1x-48D routers, how these routers fit in the Nikhef network and what the experiences of Nikhef are working with these routers.

Site Reports 2 / 20**KEK Site Report**

Authors: Go Iwai¹; Koichi Murakami¹; Ryo Yonamine¹; Soh Suzuki¹; Tomoaki Nakamura^{None}; Tomoe Kishimoto^{None}

¹ *KEK*

Corresponding Author: ryo.yonamine@kek.jp

KEK is promoting various accelerator science projects by fully utilizing the electron accelerator in Tsukuba and the proton accelerator in Tokai. These projects require a large amount of data processing, and KEKCC is operated as a central computer system to support them. In this presentation, an overview of KEKCC and its recent operation will be given. The next procurement scheduled for 2024 will also be touched upon.

Networking & Security 2 / 21**Sustainable Self-Inspection Initiatives for Improving Information Security at KEK**

Author: Ryo Yonamine¹

¹ *KEK*

Corresponding Author: ryo.yonamine@kek.jp

Public server managers at KEK are supposed to conduct a vulnerability self-inspection once a year to maintain and improve the security awareness. We have developed a new web application form dedicated to such self-inspection campaigns. The web application form also has some utility features, for instance, to generate a summary PDF file for security board meetings and JSON files for data backup. This presentation will show what we have built, what we have accomplished and challenges that remain.

Computing & Batch Services / 22**HEPiX Benchmarking Working Group Report**

Authors: Christopher Henry Hollowell¹; Domenico Giordano²; Michele Michelotto³; Randall Sobie⁴

¹ *Brookhaven National Laboratory (US)*

² *CERN*

³ *Universita e INFN, Padova (IT)*

⁴ *University of Victoria (CA)*

Corresponding Author: hollowec@bnl.gov

HEPScore is a new CPU benchmark created by the HEPiX Benchmark Working Group to replace the HEPSPEC06 benchmark that is currently used by the WLCG for procurement, computing resource pledges and performance studies.

The development of the new benchmark, based on HEP applications or workloads, has involved many contributions from software developers, data analysts, experts of the experiments, representatives of several WLCG computing centres, as well as the WLCG HEPiX Deployment Task Force.

The HEPiX benchmark has been used to show that HEP applications running on servers with ARM processors are as performant as servers with Intel and AMD processors but consume 30% less power.

This observation is a key reason for the recent work by the HEPiX Benchmark Working Group to write a plug-in for the HEPiX Suite so that the power consumption of the server can be measured during the running of the HEPiX benchmark.

In this presentation, we will report on the progress of the power measurement plug-in and present some early results.

IT Facilities & Business Continuity / 23**The new BaBar Long Term Data Analysis facility**

Author: Marcus Ebert¹

¹ *University of Victoria*

Corresponding Author: mebert@uvic.ca

If of interest, the setup of the new BaBar computing system can be described. BaBar had to move all of its computing infrastructure out of SLAC and installed a new system at UVic in 2021. While it was tried to keep the interface user are familiar with the same as before, the underlying system is using a more modern infrastructure now. It can be described how the data access is handled over WAN, how the analysis system is using Openstack VMs on demand, how the documentation had to be changed to be usable without having central manpower to manage changes, as well as how the collaboration tools (meeting system, analysis paper review, calendar, mailing lists, HN forum,...) have evolved to be usable as long as Babar is planning to do analyses. It may be interesting for any site or experiment looking ahead to plan for long term data and analysis preservation.

IT Facilities & Business Continuity / 24**Configuration management in the PDP group at Nikhef**

Author: Andrew Pickford¹

¹ *Nikhef*

Corresponding Author: andrewp@nikhef.nl

In the Nikhef PDP group we use salt to manage our systems and an extended version of reclass to store our system configuration data. The talk will cover how we do our configuration management and some lessons learnt from the way we do things.

Site Reports 3 / 25

IHEP Site Report

Authors: Xiaowei Jiang¹; 石京燕 shijy^{None}

¹ *Chinese Academy of Sciences (CN)*

Corresponding Author: xiaowei.jiang@cern.ch

The status of computing, storage, network and all related services at IHEP site

Storage and Filesystems / 26

Exploring storage technologies for HPSS disk caches

Author: Andreas Petzold¹

Co-authors: Dorin-Daniel Lobontu ; Doris Ressmann ; Preslav Konstantinov ¹

¹ *KIT - Karlsruhe Institute of Technology (DE)*

Corresponding Author: andreas.petzold@cern.ch

At KIT we operate HPSS as a tape system for the GridKa WLCG Tier-1 and for the Baden-Württemberg Data Archive service. Performance limitations of the HPSS disk cache systems led us to explore new technology options for the disk cache, based on classic storage systems with SSDs, storage servers with local NVMe devices, and also options based on IBM Storage Scale. We will present details on the different possible solutions, including benchmarks.

Site Reports 1 / 27

Canadian ATLAS Tier-1 Site Report

Author: Di Qing¹

¹ *TRIUMF*

Corresponding Author: dqing@triumf.ca

We will give a status report on the Canadian Tier-1 centre and cover several infrastructure and operational aspects, including OS plans and security initiatives.

IT Facilities & Business Continuity (C&F) / 28**Canadian ATLAS Tier-1 Analytics Infrastructure****Author:** Fernando Fernandez Galindo¹¹ TRIUMF (CA)**Corresponding Author:** ffernandezgalindo@triumf.ca

We describe the ongoing analytics project at the Canadian ATLAS Tier-1 centre whose objective is to gather, process, analyze and visualize both metrics and logs captured from the hardware and software infrastructure that build up the Tier-1 site to help monitor its health and state.

The project started in 2020 with most of the work initially focused on identifying which data to capture, how to process, store and visualize it, as well as deciding which hardware and software to utilize. We will provide a brief description of the heterogeneous nature of the data collecting infrastructure, focusing on the technologies introduced with this project: the Elasticsearch suite of tools as the main workforce for capturing, processing, and storing the data utilizing Beats, Logstash and Elasticsearch respectively; Grafana for visualization; and InfluxDB for tape library metrics. This will include a brief description of how it is set up, including example dashboards for the main datasets such as dCache, HTCondor, Linux system and security logs and tape library events.

We will also describe the hardware purchased and installed in 2022 as well as current and future work. Eventually the objective is to add machine learning methods on these datasets to provide more insights into the workings of our infrastructure, automated alerts mechanism based on predictive models, and finding correlations within the different systems to help identify sources of inefficiencies.

Storage and Filesystems / 29**Deploying and Running Ceph Clusters for Analysis Facilities at RAL****Author:** Robert Appleyard^{None}**Co-authors:** Aidan McComb ; Thomas Byrne**Corresponding Author:** rob.appleyard@stfc.ac.uk

The RAL Scientific Computing Department provides support for several large experimental facilities. These include, among others, the ISIS neutron spallation source, the Diamond X-Ray Synchrotron, the Rosalind Franklin Institute, and the RAL Central Laser Facility. We use several Ceph storage clusters to support the diverse requirements of these users.

These include Deneb, a petabyte-scale CephFS cluster, Sirius, a pure-NVMe cluster used to provide the underlying storage for STFC's private cloud, our WLCG-focussed Echo cluster which also provides S3 and SWIFT access, and Arided, a new SSD cluster providing mountable CephFS storage to our private cloud. While all of these services use Ceph to provision the storage, each has a different architecture and usage profile.

This paper will provide an outline of these services, their development and deployment, how they are used, their hardware requirements and loadings, our experiences of supporting them as production services. We will discuss the expected development roadmaps for these services for the remainder of 2023 and going into 2024, and also provide an update on recent changes to the Echo service and its XrootD interface.

Networking & Security 3 / 30

Securing the RAL campus

Author: Liam Atherton^{None}

Corresponding Author: liam.atherton@stfc.ac.uk

In the current research and education environment, the threat from cybersecurity attack is acute having grown in recent years. We must collaborate as a community to defend and protect ourselves. This requires both the use of detailed, timely and accurate threat intelligence alongside fine-grained monitoring.

We report on the development of a security operations centre for the Rutherford Appleton Laboratory to monitor both the general network and LHCOPN links. In this presentation we will share the current state of the SOC and how we aggregate, enrich and analyse the relevant data collected. We will also talk about the components of the SOC we use and how they work together to form a comprehensive system.

IT Facilities & Business Continuity / 31

Supporting Distributed Subatomic Physics Computing in Canada

Author: Leslie Groer¹

¹ *University of Toronto (CA)*

Corresponding Author: groer@physics.utoronto.ca

For over two decades, Canada has fostered and supported subatomic physics distributed computing. Many of the handful of dedicated professionals who have contributed to the establishment of these infrastructures have been involved since the early 2000s and form the heart and soul of the Digital Research Alliance of Canada's (DRAC) Subatomic Physics National Team.

The team has deployed and supported these platforms on a variety of systems. There are some unique challenges in the Canadian context as more and more consolidation on to ever larger national hosting sites has occurred where the requirements of distributed computing are not always well accommodated, especially in an era characterized by escalating apprehensions concerning research platform security.

Experience gained in these environments and tools and techniques developed in the context of supporting large and diverse research computing needs have informed the deployment of support and other infrastructures within the DRAC.

Some of this rich history, coupled with the hurdles encountered and wisdom acquired along the way, will be elaborated upon.

Computing and Batch Services / 32

Quantum Assisted Calorimeter Simulation

Authors: Abhishek Abhishek^{None}; Geoffrey Fox¹; Hao Jia²; J. Quetzalcoatl Toledo-Marin³; Maximilian J Swiatlowski⁴; Roger Melko⁵; Sebastian Gonzalez⁶; Semhimul Hoque⁷; Soren Andersen³; Tiago Vale⁸; Wojtek Fedorko³

¹ *University of Virginia*

² *University of British Columbia (CA)*

³ *TRIUMF*

⁴ TRIUMF (CA)

⁵ Perimeter Institute

⁶ UBC

⁷ University of Waterloo

⁸ Simon Fraser University (CA)

Corresponding Author: j.toledo.mx@gmail.com

Numerical simulations of collision events within the ATLAS experiment have been instrumental in shaping the design of future experiments and analyzing ongoing ones. However, the accuracy achieved in describing Large Hadron Collider (LHC) collisions comes at a substantial computational cost, with projections estimating the requirement of millions of CPU-years annually during the High Luminosity LHC (HL-LHC) run. Notably, the full simulation of a single LHC event using Geant4 currently demands approximately 1000 CPU seconds, with calorimeter simulations dominating the computational burden. Deep generative models are being developed to act as surrogates of the calorimeter data generation pipeline, and can potentially decrease the overall time to simulate single events by orders of magnitude. We introduce a novel Quantum-Assisted deep generative model. Our model combines a variational autoencoder (VAE) on the exterior with a Restricted Boltzmann Machine (RBM) in the latent space, offering enhanced expressiveness compared to conventional VAEs. RBM nodes and connections are crafted to enable the use of qubits and couplers on a D-Wave quantum annealing processor.

We will make some initial comments on the infrastructure needed for deployment at scale.

Vendor talk / 33

Sustainable Immersion Cooling for HPC: The Path to Energy Efficiency and Environmental Responsibility

Author: Eliot Ahdoot¹

¹ Hypertec Systems Inc.

Corresponding Author: eliota@hypertec.com

Join Eliot Ahdoot, Chief Innovation and Sustainability Officer at Hypertec, as he explores sustainable immersion cooling solutions in HPC. In this talk, Eliot will unveil an innovative approach that not only enhances performance but also champions environmental responsibility.

The world of scientific computing faces a persistent challenge: energy consumption. Eliot's presentation centers on a game-changing solution - Immersion Cooling. As data centers push their servers to greater power densities, a near-future surge (of at least three times) in server power requirements will overwhelm traditional air-cooling methods. With a drastic rise in global server numbers, cost reduction and the extension of infrastructure lifespan are becoming imperative.

As sustainability takes center stage in economic policies and global consciousness, every research lab and university concerned about our planet's survival and future generations' well-being must adopt environmentally friendly practices.

Key Highlights:

- **Unlocking Immersion Cooling's Potential:** Explore how immersion cooling operates and redefines efficiency in HPC environments. Real-world examples will highlight significant energy savings and the limitless scalability it offers.
- **Environmental Impact Assessment:** Delve into the eco-friendly aspects of immersion cooling, including reduced water consumption, lower carbon emissions, and a smaller environmental footprint. Discover how these solutions align with global sustainability objectives.
- **Practical Steps Toward Sustainability in HPC:** Gain actionable insights for integrating immersion cooling into your HPC infrastructure. Explore best practices, cost-effective strategies, and the roadmap to a greener future.

By the end of this presentation, you will have a comprehensive understanding of how immersion cooling can revolutionize the sustainability of HPC facilities. This innovation contributes to both scientific excellence and a more eco-conscious world. Join us in this crucial conversation, where technology meets sustainability to shape a brighter future for scientific computing.

Joint HEPiX-LHCONE session / 34**Ensuring Use of IPv6 after it is deployed****Author:** David Kelsey¹¹ *Science and Technology Facilities Council STFC (GB)***Corresponding Author:** david.kelsey@stfc.ac.uk

The HEPiX IPv6 working group has been chasing the deployment of dual-stack IPv6/IPv4 storage services in WLCG for nearly 6 years. Finally, the deployment is essentially complete with more than 97% of all LHC experiment Tier-2 storage services now IPv6-capable. There is, however, still substantial use of the legacy IPv4 protocols. The group has been identifying obstacles to the use of IPv6 and has successfully fixed many of the problems. The agreed endpoint of the IPv6 transition remains the move of all WLCG services to IPv6-only within the next few years. This talk will present all the work done and show our plans for the move to IPv6-only.

Networking & Security 3 / 35**New Security Trust and Policies - for WLCG and other Research Infrastructures****Author:** David Kelsey¹¹ *Science and Technology Facilities Council STFC (GB)***Corresponding Author:** david.kelsey@stfc.ac.uk

Many years ago, the Joint WLCG/OSG/EGEE security policy group successfully developed a suite of Security Policies for use by WLCG, EGI and others. These in turn formed the basis of the AARC Policy Development Kit, published in 2019. Many infrastructures have since used the template policies in the AARC PDK but found they had to modify them to meet their needs. The Policy Templates are now being modified, taking feedback from others into account, to new template versions in the WISE Community Security for Collaborating Infrastructures working group. In WLCG, many of the security policies are now in need of updating and revision. The work to produce new policy templates and to update WLCG security policies will be presented. This is essential for building trust within WLCG and also externally with other Infrastructures.

Joint HEPiX-LHCONE session / 36**LHCOPN and LHCONE update****Author:** Edoardo Martelli¹¹ *CERN***Corresponding Author:** edoardo.martelli@cern.ch

Latest news on LHCONE and LHCOPN development, on the on-going R&D projects, on the preparation for WLCG Data Challenge 2024

Site Reports 3 / 37

CERN site report

Author: Jarek Polok¹

¹ *CERN*

Corresponding Author: jaroslaw.polok@cern.ch

News from CERN since the last HEPiX workshop. This talk gives a general update from services in the CERN IT department.

IT Facilities & Business Continuity / 38

Oxford Computer Room Air Conditioning upgrades

Author: Peter Gronbech¹

¹ *University of Oxford (GB)*

Corresponding Author: pete.gronbech@physics.ox.ac.uk

Brief overview of upgrades to the two Oxford Computer Rooms. Discussing the improvements to PUE, but also the difficulty in ensuring that we got the improvements.

Site Reports 1 / 39

BNL Site Report

Authors: Christopher Henry Hollowell¹; Ofer Rind²; Tony Wong^{None}

¹ *Brookhaven National Laboratory (US)*

² *Brookhaven National Laboratory*

Corresponding Authors: rind@bnl.gov, tony@bnl.gov

An update on recent developments at the Scientific Data & Computing Center (SDCC) at BNL.

Computing & Batch Services / 40

The 2023 HTCondor workshop in Europe

Authors: Helge Meinhard¹; Christoph Beyer^{None}; Thomas Hartmann²

¹ *CERN*

² *Deutsches Elektronen-Synchrotron (DE)*

Corresponding Author: helge.meinhard@cern.ch

The 2023 edition of the HTCondor workshop in Europe, an annual event mostly targeted at current and future administrators of HTCondor instances, was held at IJCLab in Orsay (France, Paris region) from 19 to 22 September. This contribution will give a short report of the highlights.

Invited talk - Ocean Network Canada / 41

Ocean Networks Canada: Continuously Delivering Multidisciplinary Data from the Deep

Author: Benoit Pirene¹

¹ *Ocean Networks Canada*

Corresponding Author: bpirene@uvic.ca

Ocean Networks Canada (ONC) is one of the largest research facilities in Canada. As its name indicates, its role consists in operating and maintaining sensor networks in the ocean. This presentation will describe ONC from the perspective of its science support and societal missions with a focus on the technologies we use. The breadth of the disciplines ONC serves and the challenges with the variety of data types will be presented as our challenges. The mid-life technology upgrade path currently under consideration will be introduced with a focus on its ability to support a full scale neutrino observatory.

Site Reports 2 / 43

LHEP site report

Author: Gianfranco Sciacca¹

¹ *Universitaet Bern (CH)*

Corresponding Author: gsciacca@mail.cern.ch

The Laboratory for High Energy Physics is an institute of the Faculty of Science of the University of Bern. We present the status of the ATLAS federated Tier-2 centre and a rundown on other activities supporting physics.

Grid, Cloud & Virtualisation and Operating Systems / 44

The CERN IT Linux strategy (after the recent events in the EL eco-system)

Authors: Alex Iribarren¹; Arne Wiebalck¹

¹ *CERN*

Corresponding Author: alejandro.iribarren@cern.ch

In this presentation, we will summarise the recent events in the Enterprise Linux eco-system, starting with Red Hat's announcement to stop publicly sharing the RHEL source code and the reaction of the clone rebuilds. We will examine the impact the new situation has on the CERN use cases and the CERN IT Linux strategy.

P-One neutrino experiment invited talk / 45**The P-One ocean-based neutrino experiment - status and prospects**

Author: Steven Robertson¹

¹ *IPP / University of Alberta*

Corresponding Author: steven.robertson@ualberta.ca

The P-ONE experiment is a planned cubic-kilometer-scale neutrino telescope to be operated in the Pacific Ocean off the west coast of Vancouver Island. P-ONE will utilize infrastructure from the Oceans Network Canada (ONC) Neptune undersea cabled network to host strings of underwater optical detectors to detect light from high energy neutrino interactions in the deep ocean waters of the Cascadia basin. This presentation will summarize the physics goals, the detector design, and the status of this project with a focus on the challenges of detector controls, communication, triggering and data flow.

invited talk: digital humanities / 46**The Digital Humanities, and its Foundation for Open Social Scholarship**

Author: Ray Siemens¹

¹ *University of Victoria*

Corresponding Author: siemens@uvic.ca

The digital humanities is typically viewed as an evolving research area that exists at the intersection of computation method and the traditional pursuits of the humanities, with foundations in earlier fields aligned with inter/disciplinary evolution of computer science, dating back to the 1950s and 60s with names such as humanities informatics and humanities computing. This talk will provide an overview of the digital humanities and its typical pursuits, and focus on one of its most promising points of current research and development, the area of open, social scholarship.

New Experiments / 47**An Square Kilometer Array Regional Centre: Scaling Digital Research Infrastructure for Astronomy in Canada**

Author: Stephen Gwyn¹

¹ *National Research Council, Canadian Astronomy Data Centre*

Corresponding Author: stephen.gwyn@nrc-cnrc.gc.ca

The Square Kilometer Array (SKA) is a massive radio telescope project being built in South Africa and Australia. While observational astronomers at all wavelengths have been heavy users of high throughput computing for decades, the data rate of the SKA, 600Pb/year, far exceeds all current facilities. Earlier this year, the federal government announced that Canada will join the SKA Observatory and will provide funding for a domestic SKA Regional Centre (SRC) to support the science exploitation of the data from this facility. We will present an overview of the SKA and the international SRC Network followed by a description of the baseline plans for the Canadian SRC and how those plans fit within the context of the large astronomy projects in Canada.

Grid, Cloud & Virtualisation and Operating Systems / 49**Overview of the Coffea-Casa Analysis Facility hosted at the University of Nebraska-Lincoln****Author:** Garhan Attebury¹¹ *University of Nebraska Lincoln (US)***Corresponding Author:** garhan.attebury@cern.ch

An overview of the Coffea-Casa Analysis Facility hosted at the University of Nebraska-Lincoln. This talk will cover the technical on-prem implementation details, including networking and storage, of this Kubernetes based cluster along with the application stack supporting HEP analysis users. Integration attempts with the local USCMS Tier2 and discussion of both where we want to go and where we have failed to get will also be discussed.

Site Reports 1 / 50**The Digital Research Alliance of Canada****Author:** Patrick Mann¹¹ *Digital Alliance of Canada***Corresponding Author:** patrick.mann@alliancecan.ca

The Digital Research Alliance of Canada is a new organization, replacing the earlier organization Compute Canada, that provides compute and storage to Canadian researchers. The Alliance provides resources for the particle physics and operates the ATLAS Tier-2 facilities as well as providing compute and storage capacity for other national and international experiments. This talk provide an overview of the services and resources that currently make up the National Platform together with a brief introduction to our operational management and support practices.

Basic IT Services and End User Services 1 / 51**InvenioRDM at the SDCC****Authors:** Louis Ralph Pelosi¹; Uma Ganapathy¹**Co-authors:** Christian Lepore²; Jerome LAURET²; Michael Poat ; Robert Hancock ; Vincent Garonne¹¹ *Brookhaven National Laboratory (US)*² *Brookhaven National Laboratory***Corresponding Author:** lpelosi@bnl.gov

In contemporary digital curation landscape, the repository platform InvenioRDM stands as a potent instrument fostering scholarly communication. However, there exist inherent limitations in its platform, which require refinement to adeptly serve the diversified and dynamic scientific communities. In this presentation, we delineate a suite of pioneering extensions to InvenioRDM, orchestrated meticulously to address identified end-user requirements and to substantially enlarge its operational ambit for sPHENIX group at the Scientific Data and Computing Center (SDCC). The core objective of this presentation is to showcase how we integrated user communities based on LDAP grouping, restriction of data access through REST APIs, extension of collaborative and search capabilities, and how we leveraged the power of customizable vocabulary to tailor the tool to our customer's needs.

Through this presentation, we aim to galvanize the InvenioRDM development community towards recognizing and assimilating these vital enhancements, fostering a richer, more collaborative digital repository ecosystem for the global scientific community.

Basic IT Services and End User Services 2 / 52

The DESY User Registry

Authors: Daniel Knittel¹; Dirk Jahnke-Zumbusch^{None}

¹ DESY

Corresponding Author: dirk.jahnke-zumbusch@desy.de

DESY's service for the management of central systems with regard to accounts, groups and mailboxes – the DESY Registry – is currently in version R2. It has been fully operational for about one year.

We will touch upon useful take-aways from the migration process, look at current development (APIs, MFA integration) and other feature requests that we are evaluating for the next development cycles.

Basic IT Services and End User Services 1 / 53

Zimbra at DESY – what comes next?

Author: Dirk Jahnke-Zumbusch^{None}

Corresponding Author: dirk.jahnke-zumbusch@desy.de

DESY uses Zimbra to provide group ware services since 2014, currently version 9, and is pondering on a possible successor. Reasons, alternative products and boundary conditions as well as user requirements will be shown.

Storage and Filesystems / 54

An even more Efficient Nordic Dcache Interface to TSM

Author: Erik Mattias Wadenstein¹

¹ University of Umeå (SE)

Corresponding Author: mattias.wadenstein@cern.ch

And overview of changes coming with the Efficient Nordic Dcache Interface to TSM (ENDIT) 2.0, reasoning why, and performance plots from benchmarking and production.

Site Reports 3 / 55

FZU Site Report

Author: Jiri Chudoba¹

¹ *Czech Academy of Sciences (CZ)*

Corresponding Author: jiri.chudoba@cern.ch

The usual site report

Joint HEPiX-LHCONE session / 56

WLCG Networking Activities, Status and Plans

Author: Shawn Mc Kee¹

Co-author: Marian Babik²

¹ *University of Michigan (US)*

² *CERN*

Corresponding Author: shawn.mckee@cern.ch

We will provide a summary of the ongoing WLCG activities in networking, including perfSONAR, analytics, monitoring and the related activities of the Research Networking Technical Working Group (RNTWG): packet marking, flow labeling, transfer optimizations via tc, mtu and packet pacing.

Basic IT Services and End User Services 1 / 57

Deploying dCache in Kubernetes

Author: Tigran Mkrtchyan¹

¹ *DESY*

Corresponding Author: tigran.mkrtchyan@desy.de

The dCache project's build and test infrastructure is based on Jenkins CI and a set of virtual machines. This infrastructure is maintained by dCache developers. With the introduction of the DESY-central Gitlab server, the developers have started migrating from VM-based testing to container-based deployments. As a result, we have packaged dCache containers and Helm charts that can be used by other sites to quickly reproduce our test and build steps or to evaluate new releases on their pre-production systems, and, eventually, become a standard model of dCache deployment at the sites.

Site Reports 2 / 59

Nikhef overview and site report

Author: Dennis van Dok¹

¹ *Nikhef*

Corresponding Author: dennisvd@nikhef.nl

An overview of projects that Nikhef is involved in, and a site report with current status of the scalable computing infrastructure.

Focus is on the challenges we are facing and new ideas that are driving the direction of developments.

Site Reports 3 / 60

HIP site report

Authors: Mikael Myllymäki¹; Roope Jukkara²; Sami Lehti¹; Tomas Lindén³; Ville Salmela²

¹ *HIP*

² *CSC*

³ *Helsinki Institute of Physics (FI)*

Corresponding Author: tomas.linden@helsinki.fi

Helsinki Institute of Physics (HIP) participates in the LHC experiments ALICE, CMS and TOTEM. HIP collaborates with CSC - IT Center for Science on providing WLCG resources. The ALICE resources are part of the Nordic distributed Tier-1 resource NDGF and the CMS resources form a CMS Tier-2 called T2_FI_HIP. The HIP dCache storage was recently upgraded and the raw capacity of the new storage is 6 760 TB, which is more than three times as large as the previous system. The new storage is located about 500 km north of the location of the previous system. This site report will mainly consist of the dCache storage upgrade.

Opening Session / 61

Welcome

Author: Peter van der Reest^{None}

Co-author: Tomoaki Nakamura

Corresponding Author: peter.van.der.reest@desy.de

Opening Session / 62

Introduction to UVic and Logistics

Corresponding Author: rsobie@uvic.ca

Opening Session / 63

An introduction to HEPiX

Grid, Cloud & Virtualisation and Operating Systems / 64**Linux at DESY**

Author: Yves Kemp^{None}

Corresponding Author: peter.van.der.reest@desy.de

Grid, Cloud & Virtualisation and Operating Systems / 65**Discussion****IT Facilities & Business Continuity / 66****IDAF@DESY: Status & Outlook**

Authors: Christian Voss^{None}; Yves Kemp^{None}

Corresponding Author: christian.voss@desy.de

This presentation will go into the requirements and considerations when offering a joint analysis facility for multiple science disciplines.

Vendor talk / 67**DDN - Importance of checkpoints with LLM****Vendor talk / 68****Vendor Introduction - Weka****Basic IT Services and End User Services 2 / 69****NET2: a first example of OpenShift/OKD for Tier 2 provisioning and cluster management in US ATLAS**

Corresponding Author: eduardo.bach@cern.ch

News from the HEPiX board / 70**Election of third co-chair**

Workshop Wrap-Up & Closing Remarks / 71

Closing Remarks

Corresponding Author: peter.van.der.reest@desy.de

Workshop Wrap-Up & Closing Remarks / 72

Call for participation

Corresponding Author: m.hester@nikhef.nl